

numbers of reports and data files. Making sense of these files on a system-wide basis requires establishment of a common frame of reference. This consists of a spatial frame of reference; i.e. what happened where, as well as a temporal frame of reference, i.e. what happened when.

The spatial frame of reference comes directly from the network topology, which is generally well known to the engineers involved. Without some sort of accurate time mark in the data files, however, establishing the temporal frame of reference is challenging at best, and may be impossible. And, where possible, it often requires a great deal of effort. Indeed, in the analysis of the large blackout of August 14, 2003 affecting much of eastern North America, more time was spent ordering and manually time-tagging unsynchronized event recordings than on any other task.

This is done by looking for common, recognizable characteristics in the data files, usually by plotting them and visually inspecting the printouts. If one of the files has an accurate time tag, the other may be 'tagged' by inference. As the events get more complex, with more and more 'noise' from other things happening at the same time, this becomes more difficult and eventually, impossible. And, as the 'degrees of separation' from each record to one with an accurate time-tag increases, the expected accuracy of the time frame of that record gets worse.

For this application, extreme accuracy is not required. The events are typically examined on a cycle-by-cycle basis, with perhaps quarter-cycle resolution. Therefore, synchronization to one millisecond is adequate.

Synchrophasors

Synchrophasors are "synchronized phasor measurements," that is, measurements of ac sinusoidal quantities, synchronized in time, and expressed as phasors. With a fixed temporal reference frame, they may be used to determine useful information about operation of the grid.

For example, power flows may be monitored in real-time,

and by measuring changes in the phase shift across parts of the grid, estimates of stress and future stability can be made. Post-processing of synchrophasor data (which can also be done in real time) can extract information about low-frequency system modes, and by examining whether the amplitudes of these modes are changing, advance warning of impending instability can be given (figure 1).

Synchrophasor measurements are not practical without a widespread, economical source of accurate time. In the case of synchrophasor measurements, where one degree of system phase angle is equivalent to about 50 microseconds, synchronization at the level of 10 μs or better is desired.

System Time and Frequency Deviation

One of the oldest and most widespread uses of accurate substation time is monitoring the average system frequency. While the grid itself, and the machines connected to it, generally have a reasonable tolerance so far as small frequency offsets is concerned, many customers use the line frequency as a sort of time standard. Consider how many clocks in your own home require re-setting when the power fails, even briefly.

All of these devices use the line frequency as their time standard. The short-term frequency is less important for this sort of timekeeping than the frequency averaged over a period of hours or days. This allows the utility to control frequency during times of peak load, as one of many variables to direct power flows and ensure orderly operation of the grid. Then, when load is down (generally in the early morning hours), the frequency can be adjusted to "zero out" the accumulated time offset. System time and frequency monitors provide the information required for this, by measuring and comparing system time (from the grid frequency) to precise time.

Multi-Rate Billing

Eager to improve overall system utilization, utilities offer certain customers' incentives to use power at off-peak hours by providing them with lower energy rates during these periods. While this seems simple enough, the amounts of

money can be considerable, and a meter with a simple real-time clock that slowly drifts off (until it is reset) probably will not be acceptable as a basis for determining billing periods.

Today, the cost to provide accurate time, directly traceable to national standards, is quite reasonable. It is practical to provide an accurate clock right at the customer premises, connected directly to (or part of) the revenue meter.

Power Quality Incentives

Utilities and system operators are now considering, and in some cases implementing, financial incentives to customers to maintain power quality. Certain types of load are well-known for causing power quality impairments, such as harmonics and flicker.

These power quality impairments are measured in accordance with international standards, such as the IEC 61000-4 series. Harmonics and flicker are summarized in reporting periods, generally 10 minutes (though other periods can be used). As with multi-rate billing, the amounts of money can be significant, so customers often operate their own monitor 'in parallel' with that of the utility. Nothing causes disputes faster than monitors that disagree on the measurements, and the measurements are certain to disagree if the reporting periods are not synchronized. Providing accurate, traceable time eliminates this potential source of disagreement. Low-cost, accurate clocks are therefore an enabling technology for power quality incentives.

Using Substation Time

Navigation satellite technology has made accurate time available in the substation, generation plant, or indeed anywhere in the world with outstanding accuracy and low cost. More and more, intelligent electronic devices (IEDs) intended for use in power systems provide capability for accurate time synchronization. To implement accurate timing in your utility, you will need to choose a source of timing signals, an appropriate substation clock, and the distribution

infrastructure to provide accurate time to your IEDs.

Satellite Navigation Signals

The first and most widely used satellite global navigation system (GNS) is the US Department of Defense Global Positioning Satellite (GPS) system. This system was put into operation, in a development mode, in the 1980s and became fully operational in 1993. In addition to precise position anywhere in the world, the GPS system provides a civilian timing signal with accuracy that is now better than 10 nanoseconds.

The Russian GLONASS (Global Navigation Satellite System) provides similar capabilities to GPS. Sporadic funding of GLONASS and the resulting inconsistent satellite coverage have hampered widespread acceptance of the GLONASS system, although it is in some ways superior to GPS with respect to accuracy.

The European Space Agency (ESA) GALILEO system is the third global satellite time and navigation system to come on line. As of this writing, the GALILEO system is still in its testing phase and is not yet fully operational. As further satellites are launched, GALILEO is expected to provide similar capabilities and accuracy to the GPS and GLONASS systems. China is also considering developing its own GNS system, called Beidou.

These systems provide timing accuracy that easily exceeds the needs of the power industry. Future development in receiver technology is expected to provide the ability to receive signals from two or even more GNS systems, though existing receivers generally are limited to a single system.

Most power systems (even in Russia) now use the GPS system for accurate timing based on GNS signals. This is due to the low cost of hardware, which has been driven by widespread acceptance of GPS for consumer positioning applications.

Distributing Timing Signals

Some substation IEDs include built-in GPS receivers

Numerous applications of accurate time exist in the substation, including time tagging IED event files, synchrophasors, measuring system time deviation and multi-rate billing.

Table 1
IRIG-B
Time Code

Bit	0x	1x	2x	3x	4x	5x	6x	7x	8x	9x
x0	+000 s Reference	+100 s Minute 01	+200 s Hour 01	+300 s Day 01	+400 s Day 100	+500 s Year 01	+600 s LSP	+700 s LO 0.5	+800 s SBS 20	+900 s SBS 29
x1	+010 s Sec 01	+110 s Minute 02	+210 s Hour 02	+310 s Day 02	+410 s Day 200	+510 s Year 02	+610 s LS	+710 s TQ LSB	+810 s SBS 21	+910 s SBS 210
x2	+020 s Sec 02	+120 s Minute 04	+220 s Hour 04	+320 s Day 04	+420 s Index bit	+520 s Year 04	+620 s DSP	+720 s TQ bit 2	+820 s SBS 22	+920 s SBS 211
x3	+030 s Sec 04	+130 s Minute 08	+230 s Hour 08	+330 s Day 08	+430 s Index bit	+530 s Year 08	+630 s DST	+730 s TQ bit 3	+830 s SBS 23	+930 s SBS 212
x4	+040 s Sec 08	+140 s Index bit	+240 s Index bit	+340 s Index bit	+440 s Index bit	+540 s Index bit	+640 s LO Sign	+740 s TQ MSB	+840 s SBS 24	+940 s SBS 213
x5	+050 s Index bit	+150 s Minute 10	+250 s Hour 10	+350 s Day 10	+450 s Index bit	+550 s Year 10	+650 s LO 01	+750 s Parity	+850 s SBS 25	+950 s SBS 214
x6	+060 s Sec 10	+160 s Minute 20	+260 s Hour 20	+360 s Day 20	+460 s Index bit	+560 s Year 20	+660 s LO 02	+760 s Index bit	+860 s SBS 26	+960 s SBS 215
x7	+070 s Sec 20	+170 s Minute 40	+270 s Index bit	+370 s Day 40	+470 s Index bit	+570 s Year 40	+670 s LO 04	+770 s Index bit	+870 s SBS 27	+970 s SBS 216
x8	+080 s Sec 40	+180 s Index bit	+280 s Index bit	+380 s Day 80	+480 s Index bit	+580 s Year 80	+680 s LO 08	+780 s Index bit	+880 s SBS 28	+980 s Index bit
x9	+090 s Position P1	+190 s Position P2	+290 s Position P3	+390 s Position P4	+490 s Position P5	+590 s Position P6	+690 s Position P7	+790 s Position P8	+890 s Position P9	+990 s Position P0

Table 2 Color Coding

The colors below correspond to the background cell colors in Table 1

Framing Bits	Original IRIG Time Code Bits	Year Bits	IEEE-1344 Extensions	Index Bits
Position and reference bits, specified in IRIG Standard 200. These bits are all 8 ms wide. All other bits are either 2 or 5 ms wide.	Time of year, coded in BCD (left 4 columns) and SBS (Straight Binary Seconds), right 2 columns. Zero bits = 2 ms; One bits = 5 ms.	Coded same as time of year bits. These bits are optional in the latest version of IRIG Std. 200 and are part of the IEEE-1344 extensions.	LSP: Leap Second Pending. LS: Leap Second Direction. DSP: Summer Time Pending. DST: Summer Time Active. LO: Local Offset, +/- 15½ hours. TQ: Time Quality Bits. Parity: on all preceding data bits. For more information, see IEEE Standard 1344 or C37.118.	Bits not used for information coding. These bits are all 2 ms wide.

to provide precise timing. More commonly, however, substations include numerous IEDs which all can benefit from synchronization. Therefore, utilities have generally opted for a satellite clock in the substation, which provides timing signals to the various IEDs. The following sections describe various methods that have been used for timing distribution, with varying levels of cost, complexity and performance.

IRIG-B

The Inter-Range Instrumentation Group of the Range Commanders' Council, a US military test range body, was faced several decades ago with a problem. Each test range had developed its own, unique time code. These time codes were recorded with test data on data tape recordings. The time codes were all incompatible, which made it difficult or impossible for the ranges to exchange data. Therefore, the IRIG group set about developing a common set of time codes – the 'IRIG codes.'

These time codes have become widely used in both military and civilian applications. Particularly, the IRIG code B (IRIG-B) has become widely accepted for time distribution in substations. This time code repeats each second, and has a total of 100 bits per second. Some of these are framing (sync) bits, some are assigned for time, and some are available for control functions (tables 1 and 2). The IRIG-B code has an 'ambiguity' of one year, since it does not contain year data.

IRIG-B code may be used in either logic-level (unmodulated) format, or as an amplitude-modulated signal with a 1 kHz carrier (figure 2). The modulated IRIG signal is particularly suitable for transmission over voice-band channels, including data channels on an instrumentation tape recorder. Because of the difficulty of accurately measuring the zero crossings, modulated IRIG-B code is generally capable of an accuracy exceeding one millisecond (one period of 1 kHz), but not usually better than ten microseconds. This level of accuracy is acceptable for some but not all substation applications. Modulated IRIG inputs to IEDs are generally

transformer-isolated and provided with an automatic gain-control stage, so as to accept input signals with a wide range of amplitudes. A typical input level range is 0.1 to 10 volts peak to peak.

The unmodulated IRIG-B code can deliver accuracy limited only by the slew rate of the digital signal, much better than one microsecond and, with care, in the range of a few nanoseconds. Unmodulated IRIG-B is normally distributed at a level of 5 volts (compatible with TTL or CMOS inputs). While some IEDs couple the IRIG-B signal directly to a logic gate, best practice has used an optically-isolated input to break ground loops. Well-designed substation clocks can generally drive numerous inputs with either modulated or unmodulated IRIG-B signals, so a single clock can synchronize all the IEDs in a substation.

IEEE-1344 Extension to IRIG-B

The IRIG codes were originally developed for test range use, and a one-year ambiguity was not a limitation. Being a military code, times were always recorded in UTC (Coordinated Universal Time, military Zulu Time), so local offsets and summer time issues were also not a concern. Leap seconds (see below), which happen once or twice a year, or less, also were not a concern of the IRIG group. And the range officer would be required to guarantee that all recorders were synchronized properly before beginning a test.

Real-time operation, 24 hours per day, 365 days per year, year after year, imposes some additional requirements. The issues identified in the previous paragraph become real concerns. As part of the original Synchrophasor standard, IEEE Standard 1344-1995, an extension for the IRIG-B code was developed using the 'control bits' field to provide an additional 2 digits of year (subsequently adopted also by the IRIG standard), as well as local offset, time quality, and bits for leap second and summer time changeovers. Some IEDs support this extension, but many do not. You may find an occasional IED which requires the IEEE-1344 extension for

proper operation. This may not be well documented in the product literature.

The IEEE-1344 extension also created a new 'modified Manchester' encoding scheme, which is a digital signal having zero average value, and suitable for transmission over an optical fiber. It encodes the same 100 bit/second data stream onto a 1 kpps square-wave carrier. This modulation method has also been accepted as part of the IRIG standard.

Signal levels and connection methods for the IRIG-B code, with or without the IEEE-1344 extensions, are identical. The only difference is the use of the control bits to provide the extra information required in continuous, real-time monitoring applications (table 1). This is a firmware feature, and does not affect the hardware interface.

One Pulse per Second (1 PPS)

Many clocks can also provide a 1 PPS signal. This, like unmodulated IRIG-B, is a digital signal and the accuracy of the rising edge is a function of its slew rate. Accuracy of a few nanoseconds is typical.

However, the 1 PPS signal also has a one-second ambiguity. Each pulse is identical, and there is no way of knowing which second a pulse is associated with. Resolving this ambiguity requires a simultaneous data channel. This can be a serial channel (i.e., RS-232 data), or some other means to identify 'which second.' Because of this limitation, and the fact that the hardware requirements and performance are similar to those for unmodulated IRIG-B code, IRIG-B has generally supplanted 1 PPS for substation use. 1 PPS is still used quite commonly in standards labs, for clock comparison, because in this application the clocks can be presumed to be within a very small error of on time. The only issue is to measure this small offset, an application for which 1 PPS is very well suited.

Some IEDs do still use 1 PPS, however, and most clocks can provide this signal. Be aware of the need to resolve the one-second ambiguity. If this is done by entering the time manually from a keypad, be aware that the IED will not know

about events like leap seconds, and may not know about local time changes. Also, whenever the IED loses power, its time must be reset.

Serial (ASCII) Broadcast Time Codes

Many clocks provide a choice of ASCII strings, containing the time, which can be sent over a hardware serial port such as RS-422 or EIA-232. Generally, one of the characters in this string is transmitted 'on time,' more or less. One example of this sort of time code is IRIG Code J.

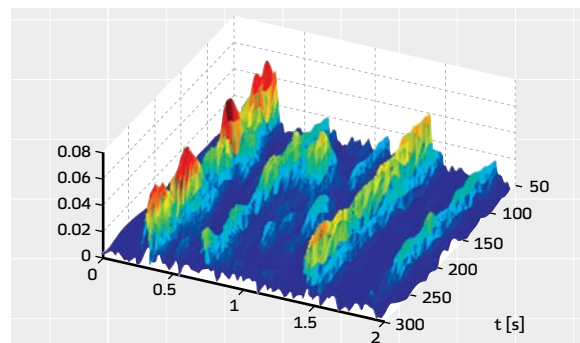
The accuracy in generating and receiving this sort of time code is a function of the hardware latency in the serial ports, usually a few bit times, plus firmware and/or software latency in the clock and client operating system. At higher bit rates, such as 19200 baud and higher, these errors can be kept below a millisecond. This sort of time code is used quite often, but primarily to drive large-digit time displays and to synchronize computers. Sometimes, the time data stream is interleaved with measurement data – when the clock also provides the function of a system time and frequency monitor. A few older IEDs were designed to accept this sort of synchronization, but this is rare today, since better methods are available.

Network Time Protocol (NTP)

NTP is a software method to transfer time between computers using a data network, such as the Internet. It provides moderate accuracy, from a few milliseconds up to a few hundred milliseconds depending on the nature of the connection between the NTP client and the server, and the performance of the computers' operating systems. NTP includes methods to estimate the round-trip path delay between the server and client, and to ignore 'outliers,' or path delay estimates which vary significantly from the typical value. However, even in the best case of a local network (e.g. Ethernet), performance is limited by the operating system stack latency, and is rarely as good as one millisecond.

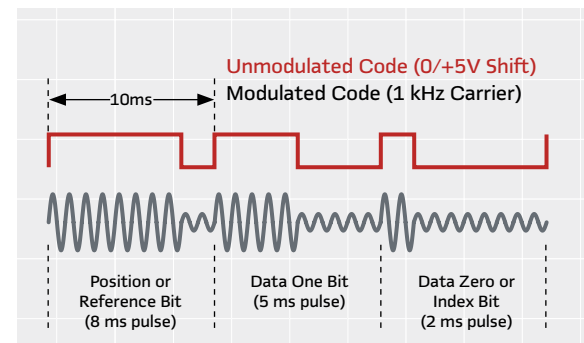
1 Phasor Measurements Graph

Results from processing of synchrophasor measurements



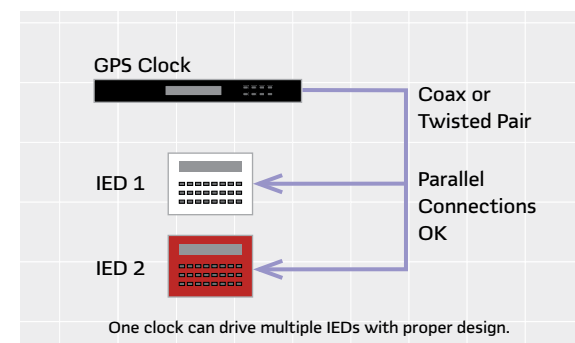
2 IRIG-B Code Structure

Modulated and unmodulated IRIG-B signals



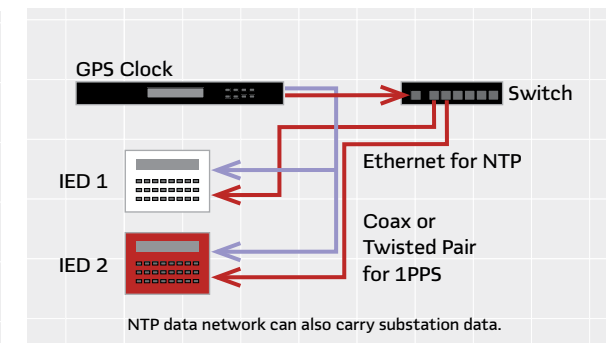
3 IRIG-B Wiring

Each IED is directly wired to the clock



4 NTP/1PPS Wiring

Requires two connections to the IED



This level of performance is adequate to resolve the one-second ambiguity of a 1 PPS signal, however, so NTP and 1 PPS together make an acceptable method of time synchronization in a substation (figure 4). This requires two connections to the IED, and so is not generally an optimal solution. We have seen in the discussion above how the dedicated logic-level input can be used alone to provide accurate time, using IRIG-B code. A natural question then is: can a network connection be used in some way to provide better time accuracy than NTP?

IEEE Standard 1588 Precision Time Protocol

This was the idea behind development of the technology described in IEEE Standard 1588: to provide hardware-level time accuracy using a standard network connection. By adding dedicated timing hardware to each port in a data network (e.g. Ethernet), the time of transmission and reception of certain messages can be determined with accuracy comparable to that of an IRIG-B or 1 PPS signal. Then, using software similar in concept to NTP, the path delays can be determined and the difference in time between a host (master) and client clock can be known (figure 5). Unlike with NTP, however, this 'hardware aided' approach can deliver timing accuracy of tens of nanoseconds.

Thus, with a single Ethernet data cable, an IED can both communicate and synchronize with sub-microsecond accuracy. This is likely to become the preferred method for both communications and synchronization in the future. Acceptance of IEEE-1588 has been slower than some may have anticipated, though, due to the effort required by vendors and users to make the system work properly (compared with IRIG-B, for example, which is quite straightforward). This consists of both firmware complexity, and the requirement for every device on a network to have the dedicated hardware capability to support the required timing measurements.

Compare this to a system using IRIG-B. With IRIG-B, all that is needed is a substation clock, synchronized to GPS or another source of accurate time. Each IED need only be

connected with a simple pair of wires, directly to the clock, and it is synchronized. With IEEE-1588, an IEEE-1588 substation clock is required. One or more switches, hubs and/or routers are also needed, and each of these must also support IEEE-1588. Finally, the IEDs must have an IEEE-1588 compatible network connection (figure 6).

In a typical substation today, many of the IEDs are still of older design and require the IRIG-B signal, and perhaps only one or two will support IEEE-1588. Obtaining and installing the necessary substation clock to support both IRIG and IEEE-1588, along with the switches and cabling for both IRIG and network signals, is not yet really cost effective. However, this will change, as vendors of substation timing and communications products are beginning to design IEEE-1588 capability into their products. In a few years, we expect that IEEE-1588 will supplant other methods of synchronization in the substation, but today, most applications still use older technology.

Time Zones, Leap Seconds and Time Scales

No discussion of substation synchronization would be complete without mentioning time zones. Early implementations of substation synchronization almost always used local time, because the operators and engineers were familiar with it. As exchanging and comparing data with neighboring utilities, over larger areas, became more common, the fact that the utilities might be in different time zones became a source of confusion. As Confucius is reputed to have said, "Man with one watch knows what time it is. Man with two watches never sure."

Coordinated Universal Time (UTC)

UTC is the modern equivalent of Greenwich Mean Time (GMT). This was originally local time at the Royal Observatory in Greenwich, England. Since modern time scales based on atomic clocks were adopted, UTC was assigned to the same time zone as GMT, to minimize confusion. Today, UTC can

be seen as a single time zone for the world, for applications where this is useful, and time tagging power system events is one of these applications. UTC can be converted to local time by adding or subtracting a fixed offset, usually but not always an integer number of hours.

More and more, utilities are changing over to UTC time, at least for communicating and storing data and event files. Modern file viewers, for example for the COMTRADE file format that is commonly used for power system event data, can read data files in UTC format and display the time to the operator in local format. In part as a result of the difficulty of comparing data from different utilities after the August 14, 2003 event in eastern North America, the North American Energy Reliability Council (NERC) is 'recommending' that utilities adopt UTC for event recording. Since these 'recommendations' have the force of law, we expect that utilities will follow them.

There are other time scales as well, which are sometimes confused with time zones. Two which you might see from time to time are International Atomic Time (TAI) and GPS time. Both of these are related to UTC and to each other by differences in leap seconds.

Leap Seconds

The rotational period of the earth is not exactly 24 hours, or more specifically 23 hours and 56 minutes, which is the nominal length of the 'sidereal' or astronomical day. As a result, the time when the celestial meridian appears overhead at any given point on the earth's surface changes slightly from day to day. This change is small, but given the high magnification of astronomical telescopes, it is important to astronomers.

UTC progresses forward at a fixed rate, based on the natural frequency of certain atomic vibrations (which define the atomic clock). As the earth's rotation diverges from this, the difference is accumulated as a correction that can be added to UTC to give sidereal time. Astronomers use this correction when directing their telescopes.

When this correction grows to approximately one second, the international timing community adds (or subtracts) one second to the UTC time scale, to keep UTC nearly aligned with sidereal time (figure 7). These added or subtracted seconds are called 'leap seconds.' By convention, they happen at midnight UTC, June 30 or December 31.

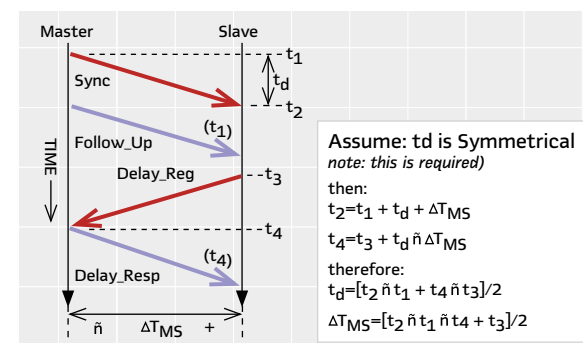
When a leap second is added, it results in two seconds having the same time in many digital time representations. The time of the leap second is reported as 23:59:60, which will when converted to binary time often have the same representation as 00:00:00 the next day. This ambiguity can cause confusion when an event occurs during or overlapping a leap second. Consequently, some people think that NERC did not go far enough in recommending UTC, but should have recommended utilities use TAI for event recording (remember that TAI could also be easily converted in a file viewer to UTC or local time).

TAI, the International Atomic Time, does not have leap seconds. As leap seconds are added to UTC, there is a difference (TAI - UTC) which is an increasing integer, presently 33 seconds. GPS time also does not have leap seconds, but it 'started' more recently than TAI, so the difference is less. Today, (GPS - UTC) is 14 seconds. And, the difference (TAI - UTC) is fixed at 19 seconds. Using TAI (or GPS) time would eliminate the potential ambiguity caused by leap seconds, but would make the 'native' time tag format less familiar to a reader.

To add a little extra complexity, there is now a discussion recommending that 'civil time,' that is, UTC, not have any more leap seconds. The difference between UTC and sidereal time would be allowed to increase past one second. If the difference ever got large enough to be a concern, then a 'leap minute' could be added to 'fix' it (this would likely happen approximately once every 100 years). If leap seconds are eliminated from UTC, the advantage of TAI would become moot, since the only difference between TAI and UTC is lack of leap seconds.

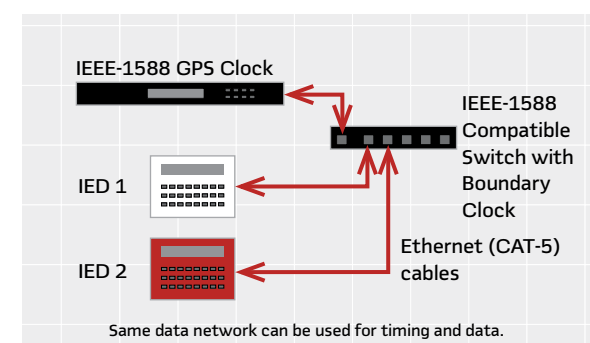
5 IEEE-1588 Synchronization

Sequence of messages for time synchronization



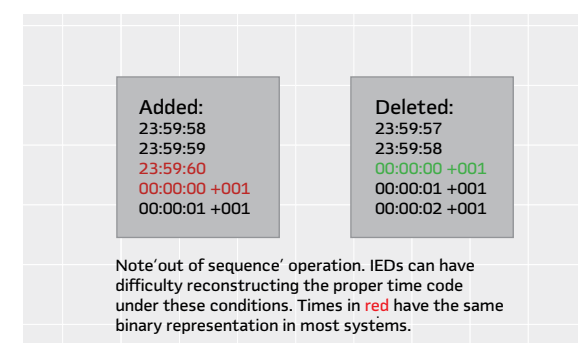
6 IEEE-1588 Wiring

An IED can communicate and synchronize over Ethernet



7 Leap Seconds

Keeps UTC nearly aligned with sidereal time



Careful consideration should be given to the choice of time format used for event recording. While local time is the obvious choice, UTC is actually better, because it reduces confusion when comparing data from different utilities in different locations.